You Can Go First—Planning For Social Law Synthesis In Asymmetric Multi-Agent Settings

Extended Abstract

Sarath Sreedharan
Colorado State University
Fort Collins, USA
sarath.sreedharan@colostate.edu

Philip Hopkins
Colorado State University
Fort Collins, USA
philip.hopkins@colostate.edu

Felipe Meneguzzi
University of Aberdeen / PUCRS
Aberdeen, UK
felipe.meneguzzi@abdn.ac.uk

ABSTRACT

Social norms are a well-known mechanism to implement agent coordination without explicit communication, runtime negotiation, or centralized planning. While synthesizing social norms in general for a multiagent system is a hard problem, we focus on a specialized problem where one of the agents accepts taking suboptimal actions, and lets the other act optimally, while avoiding potential conflicts. We show how to encode this problem into problems amenable to using existing planning systems as solvers. The resulting encodings are provably complete, guaranteeing that we can identify social norms that support such coordination.

KEYWORDS

Social Norms, Stackelberg Planning, Human-Robot Interaction

ACM Reference Format:
Sarath Sreedharan, Philip Hopkins, and Felipe Meneguzzi. 2026. You Can Go First—Planning For Social Law Synthesis In Asymmetric Multi-Agent Settings: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/FAZU7063

1 INTRODUCTION

While recent progress on autonomous agents often focuses on individual agents in isolation, acting harmoniously in heterogeneous teams involving humans remains less investigated. One mechanism to ensure coordination is social norms [11]. Recent work on automatically synthesizing such rules [7, 8, 10] mostly focuses on settings where norms are the only means of control. By relaxing this requirement, we allow novel norm formalisms.

Figure 1 shows a robot and human chef sharing a kitchen. Each chef may inadvertently sabotage the other's goal. Agents do not have equal privileges—the robot should avoid conflicts with humans. For example, rather than using oven *o1*, the robot could use *o2*. Explicitly accounting for asymmetry generates simpler social laws with less burden on the privileged agent.

We investigate synthesizing social norms for asymmetric dyadic settings. Our contributions are twofold. First, we introduce the problem of synthesizing social norms for asymmetric teams. Second, we develop Stackelberg [1] and FOND [3] planning compilations.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/FAZU7063

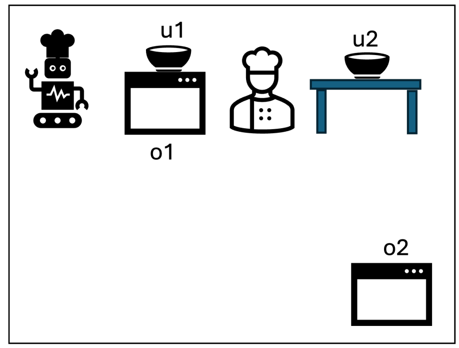

Figure 1: An example of an asymmetric multi-agent setting where potential conflicts could arise. Here, the robot chef can use either of the ovens, *o1* or *o2*, but requires the utensil *u1*. The human chef can use either the utensils or the ovens.

2 BACKGROUND AND FORMULATION

We focus on deterministic planning [6] $\mathcal{M} = \langle \mathcal{F}, \mathcal{A}, \mathcal{I}, \mathcal{G} \rangle$, where $\mathcal{F}$ is the set of propositional fluents, $\mathcal{A}$ is the set of actions (each $a = \langle Pre(a), Add(a), Del(a) \rangle$), $\mathcal{I}$ is the initial state, and $\mathcal{G}$ is the goal. A solution is a plan, i.e., a goal-achieving sequence of actions.

An asymmetric multiagent problem is $\mathcal{M}_{comb} = \langle \mathcal{M}_P, \mathcal{M}_S \rangle$, where $\mathcal{M}_P$ is the primary and $\mathcal{M}_S$ is the secondary agent model. We define our problem for a primary agent $\mathcal{P}$ and a secondary agent $\mathcal{S}$, where $\mathcal{P}$ has higher privileges. A plan π is $\mathcal{P}$-acceptable if $|\pi_{\mathcal{P}}| \leq C^*_{\mathcal{P}}$ (optimal cost for $\mathcal{P}$). We assume agents take turns to act, with $\mathcal{P}$ performing the first action.

Social norms take enforcement [4] or regimentation [5] approaches. We use regimentation, defining social norms as $\mathcal{N} = \langle \mathcal{G}^{\mathcal{N}}, \pi_{\mathcal{S}} \rangle$, where $\mathcal{G}^{\mathcal{N}} \subseteq \mathcal{F}$ are additional goal constraints for $\mathcal{P}$ and $\pi_{\mathcal{S}}$ is the exact plan for $\mathcal{S}$. This must satisfy two conditions: (C1) there exists a $\mathcal{P}$-acceptable plan achieving all goals; and (C2) no optimal $\mathcal{P}$ plan satisfying $\mathcal{G}^{\mathcal{P}} \wedge \mathcal{G}^{\mathcal{N}}$ violates the joint plan when combined with $\pi_{\mathcal{S}}$. To generate social norms efficiently, we bound the secondary agent plan length to k.

3 COMPILATIONS AND RESULTS

We cast the problem of generating social norms into finding either STACKELSAT [2] solutions or FOND [3] strong policies. The Stackelberg compilation frames the norm-synthesis problem as a leader-follower game, where the leader is the norm synthesizer and the follower is an adversarial agent attempting to invalidate the proposed norms. In this setup, the leader's objective is to identify a

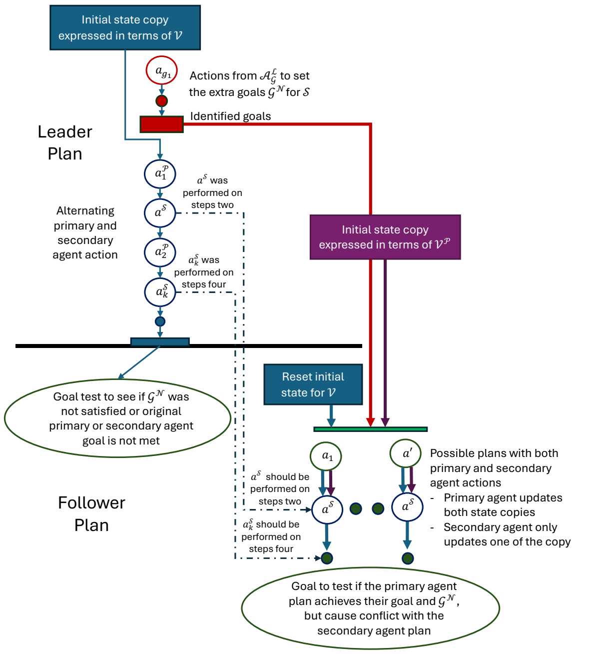

Figure 2: A visualization of the Stackelberg compilation of the norm synthesis problem.

set of additional goal constraints for the primary agent, $\mathcal{G}^{\mathcal{N}}$, and a specific plan for the secondary agent, $\pi_{\mathcal{S}}$. The leader must ensure that these proposed norms satisfy two critical conditions: (C1) there exists a primary agent-acceptable plan that achieves all goals (original and new constraints); and (C2) no optimal primary agent plan satisfying the combined goals ($\mathcal{G}^{\mathcal{P}} \wedge \mathcal{G}^{\mathcal{N}}$) conflicts with the secondary agent's plan $\pi_{\mathcal{S}}$. The follower's role is to challenge these conditions. It can invalidate the norms by demonstrating that C1 fails (i.e., no such primary agent plan exists) or by finding a conflicting primary agent plan that violates C2. To model this, the leader's actions in the planning domain include those that achieve fluents in $\mathcal{G}^{\mathcal{N}}$, execute the primary agent's actions, and also include no-op actions to pad the plan length as needed. The follower, on the other hand, has actions designed to achieve its goal of invalidation, either by showing that the resulting state does not satisfy all goals or by explicitly constructing a conflicting plan. This adversarial structure allows the Stackelberg planner to search for robust norms that hold even under an intelligent, counter-planning primary agent. Figure 2 provides a visual representation of this compilation, highlighting the interaction between the leader's norm proposal and the follower's validation/invalidation attempts.

Alternatively, we compile norm synthesis into a FOND (Fully Observable Non-Deterministic) planning instance. This approach is particularly suitable for scenarios where the primary agent's choices, while respecting the norms, can lead to different outcomes that the secondary agent must be prepared for. In a FOND compilation, we use the non-deterministic effects to model the primary agent's choices and potential behaviors within the bounds of the proposed norms. The goal of the FOND planner is then to find a strong policy for the secondary agent. A strong policy is a mapping from states to actions that guarantees goal achievement regardless of the non-deterministic outcomes (i.e., regardless of how the primary agent acts, as long as it adheres to the norms). This means the secondary agent's plan, $\pi_{\mathcal{S}}$, must be robust enough to handle any optimal or acceptable primary agent plan that satisfies the additional goal constraints $\mathcal{G}^{\mathcal{N}}$. The FOND formulation implicitly addresses condition C2 by requiring the secondary agent's policy to succeed under all possible primary agent behaviors that are consistent with the norms. If a planner cannot find a strong policy, the proposed norms are not robust, as there exists a sequence of primary agent actions (consistent with the norms) that leads to a failure state for the joint plan.

We evaluated both the Stackelberg and FOND compilation methods using state-of-the-art planners on variants of standard IPC (International Planning Competition) benchmarks. For the Stackelberg compilation, we employ a symbolic Stackelberg planner [12], which is known for its ability to handle complex leader-follower interactions. For the FOND compilation, we use the PR2 FOND planner [9], a robust planner capable of finding strong policies in non-deterministic domains. Empirical results show performance characteristics for each approach. FOND planners have superior performance on Blocksworld instances, successfully solving all test cases, including some in which the Stackelberg planner struggled or failed to find a solution. Conversely, the Stackelberg planner has better performance on some Gripper and Logistics instances, solving problems that proved challenging for the FOND planner. This suggests that for certain problem structures, the non-deterministic compilation provides an advantage over an adversarial game formulation and vice-versa. Specifically, Stackelberg planning generally struggled with larger state spaces, which is a known limitation for many symbolic planners. This highlights a crucial trade-off: while more general formulations like Stackelberg can capture complex interactions, efficient solvers tailored to specific problem structures (like FOND for strong policies) can often outperform them in their respective niches.

4 CONCLUSION

This paper formalizes social norms in asymmetric multi-agent settings and establishes the theoretical complexity of identifying such norms provided a cost bound. We develop two sound and complete methods for automatically synthesizing these norms by compiling the synthesis problem into either Stackelberg Planning or FOND planning formalisms. Our formulation has multiple practical applications, particularly in human-robot interaction settings where asymmetric planning models provide a more accurate representation of team dynamics. These social norms can serve as a foundation for automatically generating bilateral interaction protocols or contractual clauses between parties. Future work includes scenarios where the primary agent's goal is partially known or under challenging recognition of optimal plans.

ACKNOWLEDGMENTS

We thank the reviewers who engaged with our work to appreciate the merits of our paper. Sreedharan's research is supported via NSF through grant 2303019, by ONR grant number N00014-26-1-2041, and a JP Morgan Faculty Award.

REFERENCES

[1] Gregor Behnke, Benedikt Leichtmann, Pascal Bercher, Daniel Holler, Verena Nitsch, Martin Baumann, and Susanne Biundo. 2017. Help me make a dinner! Challenges when assisting humans in action planning. In *2017 international conference on companion technology (ICCT)*. IEEE, 1–6.

[2] Gregor Behnke and Marcel Steinmetz. 2024. On the Computational Complexity of Stackelberg Planning and Meta-Operator Verification. In *Proceedings of the 34th International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI Press, 20–24.

[3] Alessandro Cimatti, Marco Pistore, Marco Roveri, and Paolo Traverso. 2003. Weak, strong, and strong cyclic planning via symbolic model checking. *Artif. Intell.* 147, 1-2 (2003), 35–84.

[4] A Garcia-Camino, P Noriega, and J A Rodriguez-Aguilar. 2005. Implementing norms in electronic institutions. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, 667–673.

[5] A Garcia-Camino, C Rodriguez-Aguilar, J A nd Sierra, and W W Vasconcelos. 2009. Constraint Rule-Based Programming of Norms for Electronic Institutions. *Journal of Autonomous Agents and Multiagent Systems* 18, 1 (2009), 186–217.

[6] Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni, and Christian Muise. 2019. *An Introduction to the Planning Domain Definition Language*. Morgan & Claypool Publishers.

[7] Erez Karpas, Alexander Shleyfman, and Moshe Tennenholtz. 2017. Automated Verification of Social Law Robustness in STRIPS. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling*.

[8] Felipe Meneguzzi, Mauricio C. Magnaguagno, Munindar P. Singh, Pankaj R. Telang, and Neil Yorke-Smith. 2018. GoCo: planning expressive commitment protocols. *Autonomous Agents and Multi-agent Systems* 32, 4 (2018), 459–502.

[9] Christian Muise, Sheila A McIlraith, and J Christopher Beck. 2024. PRP Rebooted: Advancing the State of the Art in FOND Planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 20212–20221.

[10] Ronen Nir, Alexander Shleyfman, and Erez Karpas. 2020. Automated Synthesis of Social Laws in STRIPS, norm. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*. 9941–9948.

[11] Yoav Shoham and Moshe Tennenholtz. 1992. On the Synthesis of Useful Social Laws for Artificial Agent Societies (Preliminary Report), norm. In *Proceedings of the 10th National Conference on Artificial Intelligence*. 276–281.

[12] Álvaro Torralba, Patrick Speicher, Robert Künnemann, Marcel Steinmetz, and Jörg Hoffmann. 2021. Faster stackelberg planning via symbolic search and information sharing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 11998–12006.